

Amendments to the Drawings:

The attached sheets of drawings include changes to FIG. 12 to show the file system including the storage object container file 84, as disclosed in the specification on page 16 lines 17-20 and page 25 lines 5-8. The specification has been amended, as set out above, to include a new reference numeral 88 designating the file system including the storage object container file 84.

Attachment: Replacement Sheet

Annotated Sheet Showing Changes

REMARKS/ARGUMENTS

Claims 8-9, 20, 27, and 40 have been amended to more clearly distinguish the cited references.

Claims 8-9 and 20 have been amended to point out that “the client using a block level access protocol over the network to access the storage object” and “the file server accessing the storage object by accessing a file containing data of the storage object” are different from “the file server copying the file over the network” so that “the client pausing the step of writing of data to the storage object after a commit operation” is not simply a pausing of the copying as in Baweja et al. U.S. 6,564,229. Instead, claims 8-9 and 20 have been amended to recite “a second file server” and to define that the “copying” is “replicating a snapshot copy of the file from the first file server over the network to the second file server concurrent with the client using the block level access protocol over the network to write data to the storage object in the first file server; ...” Thus, with respect to applicants’ claim 8 (and similarly for claim 20), the applied references are more clearly distinguished by “the client pausing the step of writing of data to the storage object in the first file server after a commit operation, and during the pause, the client performing the step of initiating the step of the first file server replicating the snapshot copy of the file from the first file server over the network to the second file server by sending the command over the second TCP/IP connection.” Support for the amendments to claims 8-9 and 20 is found, for example, in applicants’ drawings in FIG. 1 (client 23, network 20, first file server 21, second file server 22), FIG. 6 (step 95 of “ACCESS DATA IN DATA STORAGE AREA OF THE STORAGE OBJECT CONTAINER FILE”), FIG. 7 (network block services driver 74 in client 13), FIG. 8 and 12 (network block services server 75, snapshot copy facility 76, and IP

replication facility 77 in data mover 26 [of first server 21 in FIG. 1]), FIGS. 10 and 11 (SNAP OPERATION), FIG. 12 (NBS connection 113 and SCSI OVER IP TCP connection 112), FIG. 13 and 14 (step 123 SYNC AND SUSPEND SCSI OVER IP API, SUSPEND, RESUME, step 125 CALLBACK RECEIVED, step 126 SEND SNAPSHOT OR REPLICATE COMMAND TO DATA MOVER VIA NBS TCP CONNECTION; step 128 CALLBACK RECEIVED, step 129 DLL FOR SNAPSHOT OR REPLICATION INITIATES RESUMPTION OF SCSI OVER IP OPERATION), and in applicants' specification on page 4 line 13 to page 5 line 5, page 6 lines 6-8 and 18-19, page 12 line 21 to page 13 line 2, page 13 line 10 to page 16 line 2, page 22 lines 21-24, and page 28 line 22 to page 29 line 21.

Claims 27 and 40 have been amended to more clearly distinguish the applied references (Chen et al. U.S. 7,076,509, Chen et al. 7,010,553, Rajan et al. 7,107,385, Busser US 2002/0095616) relied upon with respect to "the specification of the internal organization of the virtual direct access storage device is stored in the single file [together with data of the virtual direct access storage device]." Page 17 of the Official Action says: "It is clear from the definition of a vdisk as shown in US 7107385 ... in Fig. 5 that internal organization of the vdisk is being defined by the metadata." However, as shown in Fig. 5, "vdisk 322 is a multi-inode object comprising a special file inode [510] and at least one associated stream inode [540, 550, 560] that are managed as a single, encapsulated storage object within the file system 320." ('385 Patent col. 11 lines 39-42; see also col. 12 line 51 to col. 14 line 62.) Thus, it appears that the Official Action is construing "the specification of the internal organization of the virtual SCSI direct access storage device" so broadly as to cover the internal organization of the container file for the virtual SCSI direct access storage device. Therefore, applicants' claims 27 and 40 have

been amended to explicitly specify “a specification of an internal organization of the virtual direct access storage device for mapping of the data of the virtual direct access storage device from the single file to the data storage, ...” Support for this limitation is found in the applicants’ FIG. 5 (container file 84 including STORAGE OBJECT ATTRIBUTES 86 including INTERNAL ORGANIZATION (E.G., RAID LEVEL, STRIPING), and in applicants’ specification, page 11 lines 13-18 as follows:

Moreover, the storage object attributes may include configuration information, such as a location (bus, target and LUN) of the storage object, and an internal organization of the storage object, such as a level of redundancy in an array of disk drives (RAID level) and a striping scheme. The specified internal organization of the storage object could be used as a guide or specification for mapping of the data storage area 87 of the container file 87 to storage in the cached disk array (49 in FIG. 2).

Applicants disagree with the statement on page 19 of the Official Action suggesting that the vdisk object or LUN container file of Chen ‘509, ‘553, or Rajan ‘385 contains a specification of an internal organization of the virtual direct access storage device because col. 3, lines 41-50 of Rajan ‘385 state that the vdisk inherits the attributes of the volume it is created from. However, Rajan ‘385 col. 11 lines 29-32 further states: “In this case, the synchronous mirroring protection is not a property of the vdisk but rather a property of the underlying volume and the reliability guarantees of the file system 320.” (Emphasis added.) “Inheritance” in this context does not mean that the inherited attribute is put in the container file, but instead, “inheritance” means that that the vdisk “assumes” the underlying reliability configuration associated with the

volume upon which the file system of the vdisk is built. (See, for example, Rajan '385 col. 10 lines 50-64 et seq.) The fact that in Chen/Rajan the virtualization system 300 creates a vdisk in such a manner that is transparent to the user (Rajan '385 col. 11 lines 7-18) teaches away from including, in a container file together with data of the virtual direct access device, "a specification of an internal organization of the virtual direct access storage device for mapping of the data of the virtual direct access storage device from the single file to the data storage."

Claims 27 and 40 have also been amended to more clearly distinguish Busser US 2002/0095616 by reciting that "the attributes of the virtual SCSI direct access storage device and the data of the virtual direct access storage device are stored together in a single file in a file system ...;" Support this limitation is found in applicants' drawings in FIG. 2 (NFS 41, CIFS 42, UxFS 44), FIG. 5 (container file 84 including STORAGE OBJECT ATTRIBUTES 86 and DATA STORAGE AREA 87), FIG. 6 (steps 95 and 97), and FIG. 12 (UxFS 44, STORAGE OBJECT CONTAINER FILE 84), and in applicants' specification on page 7 lines 5-12, page 9 lines 4-11, page 16 lines 17-20, and page 25 lines 5-8. In addition, FIG. 12 of the drawings has been amended to show the file system containing the storage object container file 84, and the specification has been amended on page 16 lines 17-20, and page 25 lines 5-8 to refer this file system by a new reference number 88.

The advantages of using a container file in a file system are set out in applicants' specification on page 9 lines 4-11 as follows:

Instead of reading or writing data directly to a physical disk drive, the SCSI termination 64 reads or writes to a data storage area of the storage object 65.

The storage object, for example, is contained in a file or file system compatible with UNIX and MS-Windows. Therefore, file access protocols such as NFS and CIFS may access the storage object container file. Consequently, conventional facilities for data sharing and protection may operate upon the storage object container file. Use of a file as a container for the storage object may also exploit some file system features such as quotas, file system cache in the data mover, and block allocation on demand.

In contrast, Busser US 2002/0095616 in FIG. 3 shows stored data 90 in a disk of a RAID storage system. The stored data is divided into two categories, disk metadata 100, and customer data 200. The disk metadata 100 contains data that the controller 26 of the RAID storage system uses to assist RAID operation and management. (Busser, page 3, paragraph [0030].) The disk metadata 100 includes the RAID level 132 of the array. The RAID level 132 is the level that corresponds to the particular RAID architecture that is in use on a particular array. (Busser, page 3, paragraph [0032].) The customer data 200 is data stored on the RAID system 10 which was sent by the host computer 30. This customer data 200 may include a variety of information from the host computer 30, such as software programs, personal data, or financial data, to name a few. (Busser, page 3, paragraph [0029].)

The RAID level information in Busser paragraph [0029] appears to be on a disk header (FIG. 2) with partition information and not in a metadata file. Apparently the controller of the RAID system 10 accesses the disk metadata 100 directly, without use of a file system, and the customer data 200 may include files accessible to the host. (See Busser, page 4, paragraph [0041]; “the host and controller determines the file locations on the drives within the array.”)

Thus, it is not seen where Busser discloses that the RAID level 132 is in any file such as any of the files accessible to the host. Consequently, it is not understood how a proper combination of Busser with Chen/Rajan would result in the invention defined in applicants' amended claims 27 and 40 and their respective dependent claims 28 and 41. There is nothing in Busser that suggests modification of the teaching in Chen et al. that a logical volume is configured for a RAID level and then a file system is built on the logical volume so that the RAID level is not a property of a file or a virtual LUN.

Hashemi (US 6,934,804) does not appear to be as pertinent as Busser because Hashemi appears to disclose not much more than a general concept of striping data at the RAID level. For example, Hashemi col. 10 line 66 says:

The distribution of data and spare storage in array A300 is performed beneath the RAID layout. Each of the LUNs in the array A300 can be further partitioned or concatenated with other LUNs to form smaller or larger LUNs for defining other RAID attributes like those of striping and/or redundancy, in accordance with another embodiment of the invention.

In contrast to Busser, it is not seen where Hashemi stores a parameter used for control of striping. Instead, the striping appears to be defined by how a large LUN would be built up by concatenation of partitions of disks in the RAID array.

On page 2 paragraph 3 of the Official Action, claims 27-36 were objected to because lines 14-16 of claim 27 were redundant and claim 27 would be more clear if the word “SCSI” were deleted wherever it appears. In response, claim 27 has been amended to delete lines 14-16 and to delete the word “SCSI” wherever it appears.

On page 2 paragraph 5 of the Official Action, claims 27, 30, and 34-36 were rejected under 35 U.S.C. 103(e) as being anticipated by Chen et al. (US 7,076,509). In reply, as set out above, claims 27 and 30 have been amended to more clearly distinguish the applied references (Chen et al. U.S. 7,076,509, Chen et al. 7,010,553, Rajan et al. 7,107,385, Busser US 2002/0095616). Applicants respectfully submit that the added limitations of including, in a container file together with data of the virtual direct access device, “a specification of an internal organization of the virtual direct access storage device for mapping of the data of the virtual direct access storage device from the single file to the data storage” distinguishes Chen et al. (US 7,076,509), and the additional limitation of this single file being in a file system further distinguishes a proper combination of Chen et al. (US 7,076,509) with Busser (US 2002/0095616) and/or Hashemi (US 6934804.). Chen et al. (US 7,076,509), Busser (US 2002/0095616), and Hashemi (US 6934804) all appear to keep RAID redundancy, striping, and related information out of any container file for a virtual direct access storage device.

Chen does not suggest that the LUN file should include a specification of an internal organization for the virtual LUN such as a RAID level or a striping scheme or related information. Instead, Chen Col. 10 lines 30-42 disclose that each volume 550 is constructed from an array of physical disks 530 organized as a RAID group. Chen Col. 10 lines 43-47

teaches that within each volume may be stored one or more virtual disks (vdisks). A vdisk is a special file type in a volume that derives from a plain (regular) file, but that has associated export controls and operation restrictions that support emulation of a disk. Thus, in Chen, the RAID redundancy or striping is disclosed as a property of the volume upon which the file system is built, rather than as an attribute of the vdisk file.

As discussed above, the teaching in Chen et al. (US 7,076,509) regarding “inheritance” is not a suggestion that the RAID redundancy, striping, and related information is a property of a “vdisk” or is stored in the “vdisk.” Instead, the fact that in Chen/Rajan the virtualization system 300 creates a vdisk in such a manner that is transparent to the user (Rajan ‘385 col. 11 lines 7-18) teaches away from including, in a container file together with data of the virtual direct access device, the RAID redundancy, striping, and related information. In short, only the applicants’ novel disclosure teaches that the RAID redundancy, striping, and related information should be stored in the container file for a virtual direct access storage device, for example, so that this information is transported with the other contents of the container file to a remote file server when the container file is replicated to the remote file server.

On page 5 paragraph 7 of the Official Action, claims 7-11 and 19-23 were rejected under 35 U.S.C. 103(a) as being unpatentable over Chen et al (U.S. 7,076,509) in view of Baweja et al (U.S. 6,564,229), in further view of Bolosky et al (U.S. 2001/0052021). As discussed above, it is respectfully submitted that the amendments to the independent claims 8 and 20 more clearly distinguish Baweja et al. and therefore claims 7-11 and 18-23 are patentable under 35 U.S.C. 103(a) over Chen ‘509 in view of Baweja and Bolosky.

As amended, the independent claims 8 and 20 define a method of replication of a snapshot copy of a storage object container file from a first file server over an IP data network to a second file server concurrent with a client using a block level access protocol over the IP network for writing data to the storage object in the first file server. The client uses the block level access protocol over a first TCP/IP connection over the network to access the storage object in the first file server, and the client initiates the step of the first file server replicating the snapshot copy of the file over the network to the second file server by sending a command over a second TCP/IP connection to the first file server. Moreover, the client pauses the step of writing of data to the storage object in the first file server after a commit operation, and during the pause, the client performs the step of initiating the step of the first file server replicating the snapshot copy of the file from the first file server over the network to the second file server by sending the command over the second TCP/IP connection.

The advantage of the invention of applicants' independent claims 8 and 20 is set out in applicants' specification on page 13 line 18 to page 14 line 12 et seq.:

In the data processing system of FIG. 2, it is desired to permit the client 23 to manage backup and replication of its SCSI storage object in the data mover 26 during concurrent access to the storage object using the SCSI over IP protocol. For example, while the client 23 writes data to the data mover 26, the data mover 26 replicates the data to the second network file server 22 in FIG. 1 by transmitting a copy of the data over the IP network 20 using the NFS or CIFS protocols. One way of doing this is to provide a parallel and concurrent TCP connection between the client 23 and the data mover 26 for control of snapshot

copy and IP replication applications in the data mover 26. This method is described below with reference to FIGS. 7 to 14.

As shown in FIG. 7, the client is provided with an application program called a virtual block device manager 71 for managing backup and replication of the client's storage object 65 in the data mover 26. In order to backup or replicate a consistent view of the storage object 65, write access to the storage object by the SCSI device driver is synchronized to the backup or replication process. For example, write access of the storage object 65 is paused at the completion of a synchronous write, a commit operation for a series of asynchronous writes, or a commit of a current transaction consisting of a series of write operations. During the pause, a snapshot copy operation is initiated for the backup or replication process.

The Official Action cites Baweja for the general concept of pausing data writing and during the pause initiating the copying of the file. As introduced above, applicants' independent claims 8 and 20 have been amended so that "the client pausing the step of writing of data to the storage object after a commit operation" is not simply a pausing of the copying as in Baweja et al. U.S. 6,564,229. Baweja discloses the pausing and resuming of a data mover or copy operation in order to free system resources in order to perform other operations. (See Baweja, Abstract.)

As set out above, the applicants' invention of independent claims 8 and 20 permits a client to manage backup and replication of storage object in the data mover 26 by using snapshot copy and replication facilities upon the container file during concurrent access to the storage object using a block level access protocol. The applicants' invention enables the client to select

a consistent view of the storage object to be replicated concurrent with write access to the storage object using the block level access protocol.

The Official Action cites Bolosky for the general concept of having two concurrent TCP/IP connections, one for data transmission and one for control signaling. In particular, Bolosky discloses a control link connection for passing control information from a computer system to a media server, and a data funnel connection for passing data of multiple media from the media server to a viewer of the computer system. (Bolosky, Abstract.) The control link 18 uses the TCP/IP protocol to send commands in the form of messages between the viewer 26 and the media server controller 14, and the data funnel 20 relies upon the UDP protocol to transfer data between the viewer and the controller (although the TCP/IP protocol may be used as well). (Bolosky, page 2, paragraph [0021].) Thus, the control link connection can be used to create a data connection between the media server and the client to facilitate the exchange of data between the media server and the client at a rate substantially equal to a rate at which the client consumes data. (See Bolosky, page 9, claim 1.)


In contrast, the applicants uses a first TCP/IP connection for a block level access protocol to access the storage object in a first file server to write data to a storage object in a file in the first file server, and uses a second TCP/IP connection for sending a command to the first file server for initiating a step of the first file server replicating a snapshot copy of the file from the first file server over the network to a second file server. Thus, the two connections are for access at different levels upon the data - block level access or file level accesses. In addition, applicants' claims 8 and 20 also recite functions at the client in addition to the access at different levels in the first file server. Activity on the client is being coordinated with the concurrent

block level write and file snapshot copy and replication process on the server using the two separate TCP/IP connections.

In view of the differences set out above between the subject matter of the applicants' independent claims 8 and 20 and the applied references, it is respectfully submitted that the subject matter as a whole would not have been obvious to one of ordinary skill in the art. The problem that the inventor is trying to solve must be considered in determining whether or not the invention would have been obvious. The invention as a whole embraces the structure, properties and problems it solves. In re Wright, 848 F.2d 1216, 1219, 6 U.S.P.Q.2d 1959, 1961 (Fed. Cir. 1988). A fact finder should be aware of the distortion caused by hindsight bias and must be cautious of arguments reliant upon ex post reasoning. See KSR v. Teleflex, 550 U.S. ___ (2007), citing Graham, 383 U. S. at 36 (warning against a "temptation to read into the prior art the teachings of the invention in issue" and instructing courts to "guard against slipping into the use of hindsight.").

In view of the above, it is respectfully submitted that the application is in condition for allowance. Reconsideration is respectfully requested, and early allowance is earnestly solicited.

Respectfully submitted,



Richard C. Auchterlonie
Reg. No. 30,607

NOVAK DRUCE & QUIGG, LLP
1000 Louisiana, 53rd Floor
Houston, TX 77002
713-571-3460